Deriving a Markov Transition Model from Option-Implied Risk-Neutral Densities

Derivation Note (Implementation in Progress)

Alexander Rom

Abstract

This paper develops a framework for constructing a discrete-time Markov transition model directly from option-implied risk-neutral densities. Using the Breeden–Litzenberger result, the risk-neutral probability distribution of the underlying asset at each maturity is recovered from the cross-section of European call option prices. These marginal distributions are then discretized and embedded into a time-inhomogeneous Markov chain subject to risk-neutral martingale constraints and local transition structure. The derivation establishes existence and feasibility conditions for tri-diagonal transition matrices that preserve both marginal consistency and no-arbitrage dynamics. This work is ongoing; numerical implementation, calibration methodology, and empirical validation using market option data are the subject of ongoing development.

Standard Derivation of Breeden–Litzenberger and the Risk-neutral Density

Consider a stock price S_t and a European call option with maturity T, strike K, and time-0 price $C(K,T)$. Under the risk-neutral measure $\mathbb{Q}$, the call price is the discounted risk-neutral expectation of its payoff:

$$C(K,T) = e^{-rT}\mathbb{E}^{\mathbb{Q}}\Big[(S_T-K)^+\Big], \qquad \text{where} \quad (S_T-K)^+ = \max(S_T-K,0).$$

Let $f_Q(s)$ denote the risk-neutral probability density function of S_T, then we can write this expectation as

$$C(K,T) = e^{-rT}\int_0^\infty (s-K)^+ f_Q(s)\,ds = e^{-rT}\int_K^\infty (s-K)\,f_Q(s)\,ds.$$

First derivative with respect to the strike K. Differentiate $C(K,T)$ with respect to K. Using Leibniz's rule for differentiation under the integral sign,

$$\frac{\partial}{\partial K}\int_K^\infty (s-K)\,f_Q(s)\,ds = -(K-K)f_Q(K) \;+\; \int_K^\infty \frac{\partial}{\partial K}\Big[(s-K)f_Q(s)\Big]ds.$$

Since $(K-K)f_Q(K)=0$ and $\dfrac{\partial}{\partial K}(s-K)=-1$, we obtain

$$\frac{\partial}{\partial K}\int_K^\infty (s-K)\,f_Q(s)\,ds = \int_K^\infty (-1)\,f_Q(s)\,ds = -\int_K^\infty f_Q(s)\,ds.$$

Therefore,

$$\frac{\partial C(K,T)}{\partial K} = e^{-rT}\frac{\partial}{\partial K}\int_K^\infty (s-K)\,f_Q(s)\,ds = -e^{-rT}\int_K^\infty f_Q(s)\,ds = -e^{-rT}\mathbb{Q}(S_T\geq K).$$

Thus the *first* derivative of the call price with respect to the strike gives the *risk-neutral tail probability.*

Second derivative with respect to the strike K. Differentiate once more with respect to K:

$$\frac{\partial^2 C(K,T)}{\partial K^2} = -e^{-rT}\frac{\partial}{\partial K}\mathbb{Q}(S_T\geq K).$$

The derivative of the tail probability with respect to the threshold K is the negative of the density evaluated at K:

$$\frac{\partial}{\partial K}\mathbb{Q}(S_T\geq K) = -f_Q(K).$$

Substituting back, we obtain:

$$\frac{\partial^2 C(K,T)}{\partial K^2} = -e^{-rT}\Big(-f_Q(K)\Big) = e^{-rT}f_Q(K).$$

Solving for $f_Q(K)$, we have the Breeden–Litzenberger relation:

$$f_Q(K) = e^{rT}\frac{\partial^2 C(K,T)}{\partial K^2}.$$

Breeden–Litzenberger shows that the cross-section of call prices across strikes encodes the full risk-neutral distribution at maturity: the first derivative with respect to K gives the risk-neutral tail probability, and the second derivative (after undoing discounting) gives the risk-neutral density.

Derive Markov Chain from Risk–Neutral Densities

Setup under the Risk–Neutral Measure. Work under the risk–neutral measure $\mathbb{Q}$.

$$\text{Let} \quad 0 = t_0 < t_1 < \cdots < t_N$$

be the time grid (e.g. option maturities), and let S_t be the underlying asset price process.

$$\text{Under } \mathbb{Q}: \qquad e^{-(r-q)t}S_t \text{ is a martingale} \quad \Rightarrow \quad \mathbb{E}_{\mathbb{Q}}[S_{t_{i+1}} \mid S_{t_i}] = S_{t_i}e^{(r-q)\Delta t_i},$$

where r is the risk–free rate, q the dividend yield, and $\Delta t_i = t_{i+1} - t_i$.
From option prices (via Breeden–Litzenberger: BL), we have for each t_i a risk–neutral density

$$f_{\mathbb{Q}}(s;t_i), \quad s > 0, \qquad \text{so that for any Borel set } A \subset \mathbb{R}_+, \qquad \mathbb{Q}(S_{t_i} \in A) = \int_A f_{\mathbb{Q}}(s;t_i)\,ds.$$

Step 1: Discretize the Price Space (States). Since the BL method derives a continuous PDF, we need to convert it into a discrete one for the Markov application. So we choose a grid of price levels to be Markov states

$$s_1 < s_2 < \cdots < s_M.$$

Define bin boundaries for states as a midpoints:

$$b_{j-\frac{1}{2}} = \frac{s_{j-1} + s_j}{2}, \quad j = 2, \ldots, M,$$

and choose outer cutoffs $b_{1/2}$ and $b_{M+1/2}$ so that $[b_{1/2}, b_{M+1/2}] \supset$ most of the support of $f_{\mathbb{Q}}(\cdot;t_i)$.
For each time t_i, define the discrete marginal

$$\pi_j^{(i)} := \mathbb{Q}\big(S_{t_i} \in [b_{j-\frac{1}{2}}, b_{j+\frac{1}{2}})\big) \approx \int_{b_{j-\frac{1}{2}}}^{b_{j+\frac{1}{2}}} f_{\mathbb{Q}}(s;t_i)\,ds, \quad j = 1, \ldots, M.$$

Collect this into a row–vector

$$\pi^{(i)} = \Big(\pi_1^{(i)}, \dots, \pi_M^{(i)}\Big), \quad \pi_j^{(i)} \ge 0, \quad \sum_{j=1}^{M} \pi_j^{(i)} \approx 1.$$

Step 2: Markov Chain and General Constraints. We construct a time–inhomogeneous Markov chain

$$S_{t_0}^{\mathrm{MC}}, S_{t_1}^{\mathrm{MC}}, \dots, S_{t_N}^{\mathrm{MC}}$$

taking values in $\{s_1, \dots, s_M\}$ with transition matrices

$$P^{(i)} = \Big(P_{jk}^{(i)}\Big)_{j,k=1}^{M}, \quad P_{jk}^{(i)} = \mathbb{Q}\Big(S_{t_{i+1}}^{\mathrm{MC}} = s_k \mid S_{t_i}^{\mathrm{MC}} = s_j\Big).$$

For all transition matrices $P^{(i)}$ constructed s.t.:

∗ **Row stochasticity:**

$$\sum_{k=1}^{M} P_{jk}^{(i)} = 1, \quad P_{jk}^{(i)} \ge 0, \quad \forall j, k.$$

∗ **Marginal propagation (Markov consistency):**

$$\pi^{(i+1)} = \pi^{(i)} P^{(i)} \quad \Longleftrightarrow \quad \pi_k^{(i+1)} = \sum_{j=1}^{M} \pi_j^{(i)} P_{jk}^{(i)}, \quad k = 1, \dots, M.$$

∗ **Risk–neutral martingale constraint (row–wise):**

$$\mathbb{E}_{\mathbb{Q}}\Big[S_{t_{i+1}}^{\mathrm{MC}} \mid S_{t_i}^{\mathrm{MC}} = s_j\Big] = \sum_{k=1}^{M} P_{jk}^{(i)} s_k = s_j e^{(r-q)\Delta t_i}.$$

Step 3: Impose Tri–Diagonal (Local) Structure. For simplicity of derivation, allow transition to the 3 closest states, which imposes a local structure as follows.

From an interior state s_j ($j = 2, \dots, M-1$) the chain can only move to s_{j-1}, s_j, or s_{j+1} in one step:

$$P_{jk}^{(i)} = 0 \quad \text{if } k \notin \{j-1, j, j+1\}.$$

Define, for an interior row j: $\quad p_{j,-} := P_{j,j-1}^{(i)}, \quad p_{j,0} := P_{j,j}^{(i)}, \quad p_{j,+} := P_{j,j+1}^{(i)}.$

The key constraints for the process are:

$$p_{j,-} + p_{j,0} + p_{j,+} = 1, \tag{1}$$

$$s_{j-1} p_{j,-} + s_j p_{j,0} + s_{j+1} p_{j,+} = s_j e^{(r-q)\Delta t_i}, \tag{2}$$

$$p_{j,-},\ p_{j,0},\ p_{j,+} \ge 0. \tag{3}$$

And let $\quad \Delta s_{j,-} := s_{j-1} - s_j < 0, \quad \Delta s_{j,+} := s_{j+1} - s_j > 0, \quad m_j := s_j\Big(e^{(r-q)\Delta t_i} - 1\Big).$

From the row sum (1) we get $\quad p_{j,0} = 1 - p_{j,-} - p_{j,+}.$

Substitute $p_{j,0}$ into the martingale constraint (2):

$$\begin{aligned} s_{j-1}p_{j,-} + s_j(1-p_{j,-}-p_{j,+}) + s_{j+1}p_{j,+} &= s_j e^{(r-q)\Delta t_i}, \\ p_{j,-}(s_{j-1}-s_j) + p_{j,+}(s_{j+1}-s_j) &= s_j\Big(e^{(r-q)\Delta t_i}-1\Big). \end{aligned}$$

Equivalently, in terms of $\Delta s_{j,\pm}$ and m_j, $\quad \Delta s_{j,-}\,p_{j,-} + \Delta s_{j,+}\,p_{j,+} = m_j.$

Thus each interior row satisfies: $\begin{cases} p_{j,-}+p_{j,0}+p_{j,+}=1, \\ \Delta s_{j,-}\,p_{j,-}+\Delta s_{j,+}\,p_{j,+}=m_j, \\ p_{j,-},\, p_{j,0},\, p_{j,+}\geq 0. \end{cases}$

This is a system of 2 equations in 3 unknowns, so there we get 1 free variable.

Step 4: Parameterize the Interior Row by $p_{j,+}$. Let $x := p_{j,+}$ be the free variable.

$$\Delta s_{j,-}\,p_{j,-} + \Delta s_{j,+}\,x = m_j; \qquad \text{[solve for } p_{j,-}] \qquad p_{j,-}(x) = \frac{m_j - \Delta s_{j,+}x}{\Delta s_{j,-}}.$$

From the row sum: $\quad p_{j,0}(x) = 1 - p_{j,-}(x) - x = 1 - \frac{m_j-\Delta s_{j,+}x}{\Delta s_{j,-}} - x.$
So for an interior state transition probabilities j,

$$\boxed{p_{j,-}(x) = \frac{m_j-\Delta s_{j,+}x}{\Delta s_{j,-}}, \quad p_{j,0}(x) = 1 - \frac{m_j-\Delta s_{j,+}x}{\Delta s_{j,-}} - x, \quad p_{j,+}=x}$$

Now determine the feasible interval of x from the non–negativity constraints.

Step 5: Lower bounds from $p_{j,-}\geq 0$. Require $\quad p_{j,-}(x) = \dfrac{m_j-\Delta s_{j,+}x}{\Delta s_{j,-}} \geq 0.$

Since $\Delta s_{j,-} < 0$, multiplying both sides by $\Delta s_{j,-}$ flips the inequality:

$$m_j - \Delta s_{j,+}x \leq 0 \quad \Longleftrightarrow \quad \Delta s_{j,+}x \geq m_j.$$

Because $\Delta s_{j,+} > 0$, $\Rightarrow \quad x \geq \frac{m_j}{\Delta s_{j,+}}$. So from $p_{j,-}\geq 0$ we get a lower bound

$$x \geq L_j := \frac{m_j}{\Delta s_{j,+}}.$$

From $p_{j,+}\geq 0$ we also have the trivial bound $x \geq 0$. Combining,

$$x \geq a_j^{\text{raw}} := \max\left(0,\ \frac{m_j}{\Delta s_{j,+}}\right).$$

Step 6: Upper Bound from $p_{j,0}\geq 0$. We require: $\quad p_{j,0}(x) = 1 - \dfrac{m_j-\Delta s_{j,+}x}{\Delta s_{j,-}} - x \geq 0.$

Multiply both sides by $\Delta s_{j,-} < 0$, flipping the inequality: $\Delta s_{j,-} - \Big(m_j - \Delta s_{j,+} x\Big) - \Delta s_{j,-} x \le 0.$

Simplify the left–hand side: $\Delta s_{j,-} - m_j + \Delta s_{j,+} x - \Delta s_{j,-} x = (\Delta s_{j,+} - \Delta s_{j,-})x + (\Delta s_{j,-} - m_j).$

$$\text{Let} \qquad A_j := \Delta s_{j,+} - \Delta s_{j,-} = s_{j+1} - s_{j-1} > 0.$$

Then the inequality becomes $\qquad A_j x + (\Delta s_{j,-} - m_j) \le 0 \quad \Longleftrightarrow \quad A_j x \le m_j - \Delta s_{j,-}.$

Since $A_j > 0$, divide both sides:

$$x \le \frac{m_j - \Delta s_{j,-}}{\Delta s_{j,+} - \Delta s_{j,-}}. \qquad \text{define} \qquad b_j := \frac{m_j - \Delta s_{j,-}}{\Delta s_{j,+} - \Delta s_{j,-}}.$$

Thus from $p_{j,0} \ge 0$ we obtain the upper bound

$$x \le b_j.$$

Step 7: Feasible Interval and Choice of $p_{j,+}$. Collecting all constraints for row j,

$$p_{j,+} = x \ge 0, \qquad x \ge \frac{m_j}{\Delta s_{j,+}}, \qquad x \le \frac{m_j - \Delta s_{j,-}}{\Delta s_{j,+} - \Delta s_{j,-}}.$$

Hence the feasible interval for $p_{j,+}$ is

$$\boxed{p_{j,+} \in [a_j, b_j], \qquad a_j := \max\left(0, \frac{m_j}{\Delta s_{j,+}}\right), \qquad b_j := \frac{m_j - \Delta s_{j,-}}{\Delta s_{j,+} - \Delta s_{j,-}}.}$$

When the grid $\{s_j\}$ and step Δt_i are chosen so that $a_j \le b_j$ and $b_j \le 1$, this interval is non–empty and contained in $[0,1]$, and a valid tri–diagonal row exists.
A simple neutral choice is the midpoint

$$p_{j,+} := \frac{a_j + b_j}{2}, \Rightarrow \quad \text{then} \qquad p_{j,-} := \frac{m_j - \Delta s_{j,+} p_{j,+}}{\Delta s_{j,-}}, \qquad p_{j,0} := 1 - p_{j,-} - p_{j,+}.$$

$$\text{Satisfying our conditions:} \quad \begin{cases} p_{j,-} + p_{j,0} + p_{j,+} = 1, \\ \Delta s_{j,-}\, p_{j,-} + \Delta s_{j,+}\, p_{j,+} = m_j, \\ p_{j,-},\, p_{j,0},\, p_{j,+} \ge 0. \end{cases}$$

Step 8: Boundary Rows.
Lower Boundary $j = 1$. Allow transitions only to $\{s_1, s_2\}$:

$$p_{1,0} := P^{(i)}_{1,1}, \qquad p_{1,+} := P^{(i)}_{1,2}, \qquad p_{1,-} := 0.$$

Row sum: $\qquad p_{1,0} + p_{1,+} = 1.$

Martingale: $\quad s_1 p_{1,0} + s_2 p_{1,+} = s_1 e^{(r-q)\Delta t_i}.$

From the row sum $p_{1,0} = 1 - p_{1,+}$, substitute:

$$\begin{aligned} s_1(1 - p_{1,+}) + s_2 p_{1,+} &= s_1 e^{(r-q)\Delta t_i}, \\ s_1 + (s_2 - s_1) p_{1,+} &= s_1 e^{(r-q)\Delta t_i}. \end{aligned}$$

Let $\Delta s_{1,+} := s_2 - s_1 > 0$ and $m_1 := s_1\Big(e^{(r-q)\Delta t_i} - 1\Big)$. Then

$$\Delta s_{1,+}\, p_{1,+} = m_1 \quad \Longrightarrow \quad \boxed{p_{1,+} = \frac{m_1}{\Delta s_{1,+}}, \qquad p_{1,0} = 1 - p_{1,+}.}$$

We require $0 \le p_{1,+} \le 1$, which should hold for sensible grid/step choices.

Upper Boundary $j = M$. Allow transitions only to $\{s_{M-1}, s_M\}$:

$$p_{M,-} := P^{(i)}_{M,M-1}, \qquad p_{M,0} := P^{(i)}_{M,M}, \qquad p_{M,+} := 0.$$

Row sum: $\quad p_{M,-} + p_{M,0} = 1.$

Martingale: $\quad s_{M-1} p_{M,-} + s_M p_{M,0} = s_M e^{(r-q)\Delta t_i}.$

From $p_{M,0} = 1 - p_{M,-}$, substitute:

$$\begin{aligned} s_{M-1} p_{M,-} + s_M(1 - p_{M,-}) &= s_M e^{(r-q)\Delta t_i}, \\ s_M + (s_{M-1} - s_M) p_{M,-} &= s_M e^{(r-q)\Delta t_i}. \end{aligned}$$

Let $\Delta s_{M,-} := s_{M-1} - s_M < 0$ and $m_M := s_M\Big(e^{(r-q)\Delta t_i} - 1\Big)$. Then

$$\Delta s_{M,-}\, p_{M,-} = m_M \quad \Longrightarrow \quad \boxed{p_{M,-} = \frac{m_M}{\Delta s_{M,-}}, \qquad p_{M,0} = 1 - p_{M,-}.}$$

Again, should hold for reasonable grids and small Δt_i, these lie in $[0, 1]$.

Step 9: Full Tri–Diagonal Transition Matrix. For a fixed time step i, after computing all rows, the transition matrix is

$$P^{(i)} = \begin{pmatrix} p_{1,0} & p_{1,+} & 0 & 0 & \dots & 0 \\ p_{2,-} & p_{2,0} & p_{2,+} & 0 & \dots & 0 \\ 0 & p_{3,-} & p_{3,0} & p_{3,+} & \dots & 0 \\ \vdots & & \ddots & \ddots & \ddots & \vdots \\ 0 & \dots & 0 & p_{M-1,-} & p_{M-1,0} & p_{M-1,+} \\ 0 & \dots & 0 & 0 & p_{M,-} & p_{M,0} \end{pmatrix}.$$

Each row satisfies: $\quad \displaystyle\sum_{k=1}^{M} P^{(i)}_{jk} = 1, \qquad \sum_{k=1}^{M} P^{(i)}_{jk}\, s_k = s_j e^{(r-q)\Delta t_i}, \qquad P^{(i)}_{jk} \ge 0.$

For interior rows, the entries are given by

$$p_{j,+} \in [a_j, b_j], \qquad p_{j,-} = \frac{m_j - \Delta s_{j,+} p_{j,+}}{\Delta s_{j,-}}, \qquad p_{j,0} = 1 - p_{j,-} - p_{j,+},$$

with

$$a_j := \max\left(0,\ \frac{m_j}{\Delta s_{j,+}}\right), \qquad b_j := \frac{m_j - \Delta s_{j,-}}{\Delta s_{j,+} - \Delta s_{j,-}}.$$

For the boundaries we use the explicit 2–point formulas above for $j = 1$ and $j = M$.

Marginal Propagation and Consistency. Given $\pi^{(i)}$ and $P^{(i)}$, the Markov chain's marginal at the next time is

$$\boxed{\tilde{\pi}^{(i+1)} := \pi^{(i)} P^{(i)}.}$$

By construction, $P^{(i)}$ is row–stochastic, satisfies the risk–neutral conditional mean constraint, and is tri–diagonal (local), so $\tilde{\pi}^{(i+1)}$ is the forward–propagated discretized marginal. If the grid and time step are sufficiently fine and the BL–based $\pi^{(i)}$ are accurately approximated on the grid, then

$$\tilde{\pi}^{(i+1)} \approx \pi^{(i+1)} \quad \forall\, i.$$

Once the sequence of matrices $\{P^{(i)}\}_{i=0}^{N-1}$ is constructed, the entire option–implied term structure of risk–neutral distributions can be propagated forward on the discrete grid by repeated multiplication, $\pi^{(i+n)} \approx \pi^{(i)} P^{(i)} P^{(i+1)} \cdots P^{(i+n-1)}$. This provides a consistent, implementable discrete-time model that can be used to simulate risk–neutral price paths, to price payoffs that depend on intermediate states, and to produce transition-based quantities such as state-conditional tail probabilities and scenario transitions. The approximation $\tilde{\pi}^{(i+1)} \approx \pi^{(i+1)}$ is expected to hold when (i) the discretization error from binning the continuous density is small (fine state grid covering the mass of $f_{\mathbb{Q}}(\cdot; t_i)$), (ii) the time step Δt_i is not too large, and (iii) the local tri–diagonal structure is a reasonable discrete proxy for a diffusion-like transition between maturities. In this regime, the constructed chain preserves probability mass and matches the risk–neutral conditional mean at each state, so the forward propagation remains close to the option–implied marginals across maturities.